Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	2131	(primary adj storage) and (secondary adj storage)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:32
L2	1985485	cop\$5 or mirror\$4 or backup or back-up	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:32
L3	1221	1 and 2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:33
L4	876554	record or directory or log	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:33
L5	779	3 and 4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:33
L6	448	5 and map\$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:34
L7	9	reconcile adj command	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM TDB	OR	ON	2005/04/28 07:34
L8	2	6 and 7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:39
L9	181	6 and reset\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/04/28 07:39

L10	122	9 and switch	US-PGPUB;	OR	ON	2005/04/28 07:40
			EPO; JPO;		6	
			DERWENT;			
			IBM_TDB			
L11	4	9 and (virtual adj switch)	US-PGPUB;	OR	ON	2005/04/28 07:41
			USPAT;			
			IRM TOR			
			IBM_TDB			



Home | Login | Logout | Access Information | Alerts |

Welcome United States Patent and Trademark Office

Search Results

BROWSE

SEARCH

IEEE XPLORE GUIDE

Results for "(((primary<near/2>storage)<and>(secondary<near/2>storage)<and>(reconcile<near/2>switch)) <in>metadata)"

⊠e-mail

Your search matched 0 of 1152881 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» View Session History

» New Search

Modify Search

в Кеу

(((primary<near/2>storage)<and>(secondary<near/2>storage)<and>(reconcile<near/2

IEEE JNL IEEE Journal or

Magazine

Check to search only within this results set

IEE JNL

IEE Journal or Magazine

Display Format:

© Citation C Citation & Abstract

IEEE CNE

IEEE Conference Proceeding

IEE Conference

No results were found.

IEE CNF

STD

BEEE

Proceeding **IEEE Standard**

Please edit your search criteria and try again. Refer to the Help pages if you need assistance revisir

Help Contact Us Privacy &:

© Copyright 2005 IEEE --

indexed by # Inspec Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

writing first data to primary storage and second data to second



THE AUS DIG TALL BRAK

Feedback Report a problem Satisfaction survey

Terms used writing first data to primary storage and second data to secondary storage

Found **94,920** of **154,226**

Sort results

relevance by Display expanded form results

Save results to a Binder ? Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 200

window

Result page: **1** $\frac{2}{3}$ $\frac{3}{4}$ $\frac{5}{5}$ $\frac{6}{7}$ $\frac{7}{8}$ $\frac{9}{9}$ $\frac{10}{10}$

Relevance scale 🔲 📟 🚾 🛭

Best 200 shown RAID: high-performance, reliable secondary storage

Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson June 1994 ACM Computing Surveys (CSUR), Volume 26 Issue 2

Full text available: pdf(3.60 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to improve aggregate I/O performance. Today they appear in the product lines of most major computer manufacturers. This article gives a comprehensive overview of disk arrays and provides a framework in which to organize current and future work, First, the article introduces disk technology and reviews the driving forces that have popularized disk arrays; performance and reliability. It discusses the tw ...

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

Integrated document caching and prefetching in storage hierarchies based on Markovchain predictions

Achim Kraiss, Gerhard Weikum

August 1998 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 7 Issue 3

Full text available: pdf(603.01 KB) Additional Information: full citation, abstract, citings, index terms

Large multimedia document archives may hold a major fraction of their data in tertiary storage libraries for cost reasons. This paper develops an integrated approach to the vertical data migration between the tertiary, secondary, and primary storage in that it reconciles speculative prefetching, to mask the high latency of the tertiary storage, with the replacement policy of the document caches at the secondary and primary storage level, and also considers the interaction of these policies with ...

Keywords: Caching, Markov chains, Performance, Prefetching, Scheduling, Stochastic modeling, Tertiary storage

Comparison of access methods for time-evolving data Betty Salzberg, Vassilis J. Tsotras June 1999 ACM Computing Surveys (CSUR), Volume 31 Issue 2



Full text available: pdf(529.53 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

Keywords: I/O performance, access methods, structures, temporal databases

Fortran 8X draft

Loren P. Meissner

December 1989 ACM SIGPLAN Fortran Forum, Volume 8 Issue 4

Full text available: pdf(21.36 MB) Additional Information: full citation, abstract, index terms

Standard Programming Language Fortran. This standard specifies the form and establishes the interpretation of programs expressed in the Fortran language. It consists of the specification of the language Fortran. No subsets are specified in this standard. The previous standard, commonly known as "FORTRAN 77", is entirely contained within this standard, known as "Fortran 8x". Therefore, any standard-conforming FORTRAN 77 program is standard conforming under this standard. New features can b ...

External memory algorithms and data structures; dealing with massive data Jeffrey Scott Vitter

June 2001 ACM Computing Surveys (CSUR), Volume 33 Issue 2

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(828.46 KB) terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

File and storage systems: The Google file system Sanjay Ghemawat, Howard Gobioff, Shun-Tak Leung

October 2003 Proceedings of the nineteenth ACM symposium on Operating systems principles

Full text available: pdf(275.54 KB) Additional Information: full citation, references, index terms

Keywords: clustered storage, data storage, fault tolerance, scalability

OceanStore: an architecture for global-scale persistent storage John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Chris Wells, Ben Zhao



November 2000 Proceedings of the ninth international conference on Architectural support for programming languages and operating systems, Volume 28, 34 Issue 5, 5

Full text available: pdf(166.53 KB)

Additional Information: full citation, abstract, references, citings, index terms

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers. data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks; monitoring also enhances performance through pro-active movement ...

OceanStore: an architecture for global-scale persistent storage John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishan Gummadi, Sean Rhea, Hakim Weatherspoon, Wes'tley Weimer, Chris Wells, Ben



Full text available: pdf(1.47 MB)

Additional Information: full citation, abstract, references, index terms

OceanStore is a utility infrastructure designed to span the globe and provide continuous access to persistent information. Since this infrastructure is comprised of untrusted servers, data is protected through redundancy and cryptographic techniques. To improve performance, data is allowed to be cached anywhere, anytime. Additionally, monitoring of usage patterns allows adaptation to regional outages and denial of service attacks: monitoring also enhances performance through pro-active movement ...

File servers for network-based distributed systems

Liba Svobodova

December 1984 ACM Computing Surveys (CSUR), Volume 16 Issue 4

Full text available: pdf(4.23 MB) Additional Information: full citation, references, citings, index terms, review

10 Consistent and automatic replica regeneration

Haifeng Yu, Amin Vahdat

February 2005 ACM Transactions on Storage (TOS), Volume 1 Issue 1

Full text available: pdf(372.24 KB) Additional Information: full citation, abstract, references, index terms

Reducing management costs and improving the availability of large-scale distributed systems require automatic replica regeneration, that is, creating new replicas in response to replica failures. A major challenge to regeneration is maintaining consistency when the replica group changes. Doing so is particularly difficult across the wide area where failure detection is complicated by network congestion and node overload. In this context, this article presents Om, the first read/write peer- ...

Keywords: Peer-to-peer storage systems, availability, consistency, regeneration, replication

11 The HP AutoRAID hierarchical storage system

John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan

February 1996 ACM Transactions on Computer Systems (TOCS), Volume 14 Issue 1

Full text available: pdf(1.82 MB)

Additional Information: full citation, abstract, references, citings, index terms

Configuring redundant disk arrays is a black art. To configure an array properly, a system



administrator must understand the details of both the array and the workload it will support. Incorrect understanding of either, or changes in the workload over time, can lead to poor performance. We present a solution to this problem: a two-level storage hierarchy implemented inside a single disk-array controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

12 <u>Distributed, object-based programming systems</u>

Roger S. Chin, Samuel T. Chanson

March 1991 ACM Computing Surveys (CSUR), Volume 23 Issue 1

Full text available: pdf(2.97 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The development of distributed operating systems and object-based programming languages makes possible an environment in which programs consisting of a set of interacting modules, or objects, may execute concurrently on a collection of loosely coupled processors. An object-based programming language encourages a methodology for designing and creating a program as a set of autonomous components, whereas a distributed operating system permits a collection of workstations or personal computers ...

Keywords: capability scheme, distributed operating systems, error recovery, method invocation, nested transaction, object model, object reliability, object-based programming languages, processor allocation, resource management, synchronization, transaction

13 Access methods for multiversion data

David Lomet, Betty Salzberg

June 1989 ACM SIGMOD Record, Proceedings of the 1989 ACM SIGMOD international conference on Management of data, Volume 18 Issue 2

Full text available: pdf(1.11 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

We present an access method designed to provide a single integrated index structure for a versioned timestamped database with a non-deletion policy. Historical data (superceded versions) is stored separately from current data. Our access method is called the Time-Split B-tree. It is an index structure based on Malcolm Easton's Write Once B-tree. The Write Once B-tree was developed for data stored entirely on a Write-Once Read-Many or WORM optical ...

14 The HP AutoRAID hierarchical storage system

J. Wilkes, R. Golding, C. Staelin, T. Sullivan

December 1995 ACM SIGOPS Operating Systems Review, Proceedings of the fifteenth ACM symposium on Operating systems principles, Volume 29 Issue 5

Full text available: pdf(1.60 MB)

Additional Information: full citation, references, citings, index terms

15 Distributed file systems: concepts and examples

Eliezer Levy, Abraham Silberschatz

December 1990 ACM Computing Surveys (CSUR), Volume 22 Issue 4

Full text available: pdf(5.33 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical





configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

16 Efficient suffix trees on secondary storage

David R. Clark, J. Ian Munro

January 1996 Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms

Full text available: pdf(1.03 MB)

Additional Information: full citation, references, citings, index terms

17 Petal: distributed virtual disks

Edward K. Lee, Chandramohan A. Thekkath

September 1996 Proceedings of the seventh international conference on Architectural support for programming languages and operating systems, Volume 31, 30 Issue 9, 5

Full text available: pdf(1.10 MB)

Additional Information: full citation, abstract, references, citings, index terms

The ideal storage system is globally accessible, always available, provides unlimited performance and capacity for a large number of clients, and requires no management. This paper describes the design, implementation, and performance of Petal, a system that attempts to approximate this ideal in practice through a novel combination of features. Petal consists of a collection of network-connected servers that cooperatively manage a pool of physical disks. To a Petal client, this collection appear ...

18 Concepts and capabilities of a database computer

Jayanta Banerjee, David K. Hsiao, Richard I. Baum

December 1978 ACM Transactions on Database Systems (TODS), Volume 3 Issue 4

Full text available: pdf(2.79 MB)

Additional Information: full citation, abstract, references, citings, index terms

The concepts and capabilities of a database computer (DBC) are given in this paper. The proposed design overcomes many of the traditional problems of database system software and is one of the first to describe a complete data-secure computer capable of handling large databases. This paper begins by characterizing the major problems facing today's database system designers. These problems are intrinsically related to the nature of conventional hardware and can only be solved by i ...

Keywords: clustering, content-addressable memory, database computers, keywords, mass memory, performance, security, structure memory

19 Proceedings of the SIGNUM conference on the programming environment for development of numerical software

March 1979 ACM SIGNUM Newsletter, Volume 14 Issue 1

Full text available: pdf(5.02 MB) Additional Information: full citation

²⁰ Private information storage (extended abstract)

Rafail Ostrovsky, Victor Shoup

May 1997 Proceedings of the twenty-ninth annual ACM symposium on Theory of computing



Full text available: pdf(1.33 MB) Additional Information: full citation, references, citings, index terms

Results 1 - 20 of 200 Result page: 1 $\underline{2}$ $\underline{3}$ $\underline{4}$ $\underline{5}$ $\underline{6}$ $\underline{7}$ $\underline{8}$ $\underline{9}$ $\underline{10}$ \underline{next}

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player